

REPORT DOCUMENTATION PAGE

Public reporting burden for this collection of information is estimated to average 1 hour per response, including reviewing the data needed, and completing and reviewing the collection of information. Send comments and suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0635).

AFRL-SR-BL-TR-00-

gathering
collection of
way, Suite

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE		3. R	
4. TITLE AND SUBTITLE Search Control for Automatic Plan Generation				5. FUNDING NUMBERS F49620-96-1-0403	
6. AUTHOR(S) Dr Martha Pollack					
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Department of Computer Science University of Pittsburgh Pittsburgh, PA 15260				8. PERFORMING ORGANIZATION REPORT NUMBER N/A	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES N/A					
12a. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release; Distribution unlimited.				12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) In this project, we analyzed the effectiveness of alternative search control strategies for automatic plan generation. We completed a detailed empirical study of control strategies for partial-order causal-link planners, identifying the relevant advantages and disadvantages of both strategies in the literative and novel strategies. We also developed search control heuristics for conditional planning, and demonstrated their effectiveness. Finally, we developed techniques for generating monitors during planning, without incurring excessive overhead. 20001127 023					
14. SUBJECT TERMS Planning, Heuristic Search, Search Control				15. NUMBER OF PAGES	
				16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL		

*logged 1-17
98/10/8*

Final Technical Report

Principal Investigator:	Dr. Martha E. Pollack
Institution:	Department of Computer Science University of Pittsburgh Pittsburgh, PA 15260
Email:	pollack@cs.pitt.edu
Grant Number:	F49620-96-1-0403
Title:	Search Control for Automatic Plan Generation
Prepared for:	Dr. Abraham Waksman Air Force Office of Scientific Research AFOSR, Bolling Air Force Base Washington, D.C. 20332-6448
Period Covered:	Aug. 1, 1996 - Jan. 31, 1998

1 Introduction

This document summarizes the research conducted on the AFOSR-supported project "Search Control for Automatic Plan Generation", Contract F49620-96-1-0403 during the period between Aug. 1, 1996 and Jan. 31, 1998 (18 months). The goal of the project has been to analyze the effectiveness of alternative search control strategies for automatic plan generation and to investigate the interactions between search control strategies and other aspects of the planning architecture. Our efforts were focused on developing strategies for:

- search control in partial-order causal-link planning;
- search control in conditional planning;
- monitor-establishment in dynamic planning, with an emphasis on the interaction between monitoring and the efficiency of planning.

In addition, early in the project we completed some relevant work that had been begun prior to the project start-date, involving search control for planners operating in domains in which actions have explicit costs associated with them.

This report is organized around these topics. We briefly describe the work we did on each of these topics, followed by a list of project-sponsored publications, which provide more details of the work. Copies of these publications are included as an appendix to this report.

2 Search Control for Partial-Order Causal Link Planning

Much of the current research in plan generation centers on partial-order causal link (POCL) algorithms, which descend from McAllester and Rosenblitt's [9] SNLP algorithm. POCL planning involves searching through a space of partial plans, where the successors of a node representing partial plan P are refinements of P . As with any search problem, POCL planning requires effective search control strategies.

In POCL planning, search control has two components. The first, *node selection*, involves choosing which partial plan to refine next. Once a partial

plan has been selected for refinement, the planner must then perform *flaw selection*, which involves choosing either a threat to resolve or an open condition to establish.

Over the past few years, several studies have compared the relative efficiency of alternative flaw selection strategies for POCL planning and their extensions [11, 8, 13, 6, 18]. These studies have been motivated at least in part by a tension between the attractive formal properties of the POCL algorithms, and the limitations in putting them to practical use that result from their relatively poor performance. To date, the POCL algorithms cannot match the efficiency of the so-called industrial-strength planners such as SIPE [16, 17] and O-Plan [4, 14]. Flaw selection strategy has been shown to have a significant effect on the efficiency of POCL planning algorithms, and thus researchers have viewed the design of improved flaw selection strategies as one means of making POCL planning algorithms more practical.

In the current project, we completed an extensive experimental study of the relative performance of the main control strategies that have been proposed in the prior literature for partial-order causal-link planning. Our results are presented in [12], in which we review the literature on flaw selection strategies, and present new experimental results that generalize the earlier work and explain some of the discrepancies in it. In particular, we describe the Least-Cost Flaw Repair (LCFR) strategy developed and analyzed by Joslin and Pollack [8], and compare it with other strategies, including Gerevini and Schubert's ZLIFO strategy [6]. LCFR and ZLIFO make very different, and apparently conflicting claims about the most effective way to reduce search-space size in POCL planning. We resolve this conflict, arguing that much of the benefit that Gerevini and Schubert ascribe to the LIFO component of their ZLIFO strategy is better attributed to other causes.

More specifically, we showed that neither the LCFR nor ZLIFO flaw selection strategy consistently generates smaller search spaces, but that by combining LCFR's least-cost approach with the delay of separable threats that is included in the ZLIFO strategy, we obtain a strategy—LCFR-DSep—whose space performance was nearly always as good as the better of LCFR or ZLIFO on a given problem. We therefore concluded that much of ZLIFO's advantage relative to LCFR is due to its delay of separable threats rather than to its use of a LIFO strategy. Although we were unable to resolve the question of whether least-cost selection is required for unforced, as well as forced flaws, we found no evidence that a LIFO strategy for unforced flaws was better. On

the other hand, separable-threat delay is clearly advantageous.

We also considered the question of computation time, and showed that often LCFR-DSep only requires computation time comparable to that of ZLIFO. LCFR-DSep can therefore be seen as paying for its own computational overhead by its search-space reduction.

Those conclusions, however, are tempered by the fact that for certain clusters of problems, our combined strategy, LCFR-DSep, does not generate minimal search spaces. In sum, as a result of our experiments we now understand the reasons that LCFR and ZLIFO perform the way they do, and how to combine the best features of both to create good default strategies for POCL planning. At the same time, it is clear that certain domain-dependent characteristics such as those we identified in several of the domains we studied must still be taken into account in settling on a flaw selection strategy for any particular planning domain.

3 Search Control for Conditional Planning

Conditional planning is an important extension to traditional planning. Conditional planners allow for conditional actions with multiple possible outcomes and for sensing actions that allow agents to determine the current state [1, 5, 7, 3]. A key question in conditional planning is: how many, and which of the possible execution failures should be planned for? One cannot, in general, plan for all the failures that can be anticipated: there are simply too many. But neither can one ignore all the possible failures, or one will fail to produce sufficiently flexible plans. Essentially, this question can be viewed as one of search control: which portion of the plan space should be searched first, to provide the highest-quality contingency plans?

In the current project, we developed Mahinur, a probabilistic partial-order planner that supports conditional planning with contingency selection; our work on this is reported in [10]. Mahinur implements an iterative refinement planning algorithm that identifies the contingencies that contribute the most to the plan's overall value, and gives priority to the contingencies whose failure would have the greatest negative impact. We concentrated on two aspects of the problem, namely, planning methods for an iterative conditional planner and a method for computing the negative impact of possible sources of failure.

We conducted experiments with reasoning about the first implementation

of Mahinur, and compared its performance to other probabilistic conditional planners, notably the C-Buridan [5] system, the best-known alternative for partial-order contingency planning. Mahinur differs from C-Buridan and other earlier systems in explicitly calculating the expected value of handling alternative contingencies at plan time. Our experiments showed that these calculations result in a significant increase in Mahinur's planning efficiency, relative to C-Buridan. We are continuing to work on the Mahinur system, incorporating methods for reducing the probability of failure by adding more supporting actions, and implementing a much larger real-world domain to use as the basis of extended experimental analyses.

4 Monitor-Selection in Planning

A further extension to planners is required when agents are situated in dynamic environments. There, a central challenge is to be appropriately sensitive to changes in its environment. In general, it is too costly to be responsive to every environmental feature that the agent knows about. On the other hand, an agent that is completely unresponsive may fail to take advantage of circumstances that can improve its plans and/or shorten its planning time considerably. The need to balance sensitivity to environmental change against appropriate stability of the plans being formed is strongly reminiscent of the ideas that led to the design of the IRMA architecture and filtering strategy in our earlier work [2].

In recent work on this project, we have introduced the idea of *rationale-based monitoring*, reported in [15]. In this approach, planning is strongly identified as a decision making process and the planning system records the rationale for the choices it makes. Even when planning consists mainly in task decomposition, it will typically involve choosing between alternatives, and the reasons for those choices constitute the plan rationale. The agent can then focus its attention on those changes in the environment that would affect the truth-value of the planning rationale.

A novel aspect of our approach is that we not only monitor features of the world that affect the current plan, but also features of the world that played a role in the decision to select that plan over alternative possibilities. We maintain two sets of monitors: plan-based and alternative-based. Every time the agent needs to make a decision among alternatives, it deliberates and selects

a particular plan. The selected plan gives rise to the plan-based monitors. At the same time, the alternatives considered give rise to alternative-based monitors. As the world state is dynamically changing, the agent remembers alternatives that it judged less valuable, monitoring the world state to see if that judgement should be changed. We implemented a prototype version of rationale-based monitoring and conducted preliminary experiments showing that it can lead to improved plans without significant overhead.

5 Project-Supported Publications¹

1. Eithan Ephrati, Martha E. Pollack, and Marina Milshtein, "A Cost Directed Planner: Preliminary Report," in *Proceedings of the 13th National Conference on Artificial Intelligence*, pp. 1223-1228, 1996.
2. David Joslin and Martha E. Pollack, "Is 'Early Commitment' in Plan Generation Ever a Good Idea," in *Proceedings of the 13th National Conference on Artificial Intelligence*, pp. 1188-1193, 1996.
3. Martha E. Pollack, David Joslin, and Massimo Paolucci, "Flaw-Selection Strategies for Partial-Order Planning," *Journal of Artificial Intelligence Research*, 6:233-262, 1997.
4. Nilufer Onder and Martha E. Pollack, "Contingency Selection in Plan Generation," in *Proceedings of the Fourth European Conference on Planning*, Toulouse, France, Sept. 1997.
5. Nilufer Onder and Martha E. Pollack, "Contingency Selection in Plan Generation," in *1996 AAAI Fall Symposium on Plan Execution*, Boston, MA, Nov. 1996. (This is a preliminary version of the previous reference.)
6. Nilufer Onder and Martha E. Pollack, "Handling Contingency Selection Using Goal Values," in *1997 AAAI Workshop on Abstractions, Decisions, and Uncertainty*, Providence, RI, July, 1997.
7. Yagil Ronen, Daniel Mosse', and Martha E. Pollack "Value-Density Algorithms for the Deliberation Scheduling Problem," to appear in *IEEE Expert*, 1998.
8. Manuela M. Veloso, Martha E. Pollack, and Michael T. Cox, "Rationale-Based Monitoring for Continuous Planning in Dynamic Environments" to appear in *Proceedings of the Fourth International Conference on AI Planning Systems*, June, 1998.

¹The first two papers listed were completed prior to the start date of the current contract, and thus do not acknowledge this contract. However, they are both within the scope of the current effort, and follow-on work, which was reported when the papers were presented at the conference, was done during the current contract period.

References

- [1] Jim Blythe. *Planning under uncertainty in dynamic domains*. PhD thesis, Carnegie Mellon University, 1997.
- [2] Michael E. Bratman, David J. Israel, and Martha E. Pollack. Plans and resource-bounded practical reasoning. *Computational Intelligence*, 4:349-355, 1988.
- [3] Greg Collins and Louise Pryor. Planning under uncertainty: some key issues. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, 1995.
- [4] Ken Currie and Austin Tate. O-plan: The open planning architecture. *Artificial Intelligence*, 52:49-86, 1991.
- [5] Denise Draper, Steve Hanks, and Daniel Weld. Probabilistic planning with information gathering. In *Proceedings of the 2nd International Conference on Artificial Intelligence Planning Systems*, pages 31-36, 1994.
- [6] Alfonso Gerevini and Lenhart Schubert. Accelerating partial-order planners: Some techniques for effective search control and pruning. *Journal of Artificial Intelligence Research*, 5:95-137, 1996.
- [7] Robert Goldman and Mark Boddy. Epsilon-safe planning. In *Proceedings of the Tenth International Conference on Uncertainty in Artificial Intelligence*, 1994.
- [8] David Joslin and Martha E. Pollack. Least-cost flaw repair: A plan refinement strategy for partial-order planning. In *Proceedings of the Twelfth National Conference on Artificial Intelligence (AAAI)*, pages 1004-1009, Seattle, WA, 1994.
- [9] David McAllester and David Rosenblitt. Systematic nonlinear planning. In *Proceedings of the Ninth National Conference on Artificial Intelligence*, pages 634-639, Anaheim, CA, 1991.
- [10] Nilufer Onder and Martha E. Pollack. Contingency selection in plan generation. In *Proceedings of the 4th European Conference on Planning*, 1997.

- [11] Mark Peot and David E. Smith. Threat-removal strategies for partial-order planning. In *Proceedings of the Eleventh National Conference on Artificial Intelligence*, pages 492-499, Washington, D.C., 1993.
- [12] Martha E. Pollack, David Joslin, and Massimo Paolucci. Flaw selection strategies for partial-order planning. *Journal of Artificial Intelligence Research*, 6:223-262, 1997.
- [13] Raghavan Srinivasan and Adele E. Howe. Comparison of methods for improving search efficiency in a partial-order planner. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, pages 1620-1626, 1995.
- [14] A. Tate, B. Drabble, and J. Dalton. Reasoning with constraints within o-plan2. Technical Report ARPA-RL/O-Plan2/TP/6 V. 1, AIAI, Edinburgh, 1994.
- [15] Manuela M. Veloso, Martha E. Pollack, and Michael T. Cox. Rationale-based monitoring for continuous planning in dynamic environments. To appear in *Proceedings of the Fourth International Conference on AI Planning Systems (AIPS-98)*.
- [16] David E. Wilkins. *Practical Planning: Extending the Classical AI Paradigm*. Morgan Kaufmann, San Mateo, CA, 1988.
- [17] David E. Wilkins and Robert V. Desimone. Applying an AI planner to military operations planning. In Mark Fox and Monte Zweben, editors, *Intelligent Scheduling*, pages 685-708. Morgan Kaufmann Publishers, San Mateo, CA, 1994.
- [18] Mike Williamson and Steve Hanks. Flaw selection strategies for value-directed planning. In *Proceedings of the Third International Conference on Artificial Intelligence Planning Systems*, pages 237-244, 1996.